INTERNAL — RESTRICTED

Microsoft 365 E5 Architecture

and the Copilot (Anthropic) Overlay

A Regulatory Obligations Reference for a Global Systemically Important Financial Institution

Version 2.2

Baseline Date: 2026-05-05

Audience: Enterprise Architects, Identity & AiRB Leads, Legal, Compliance & Risk Officers

Document Control

Field	Value
Title	M365 E5 Architecture and the Copilot (Anthropic) Overlay — Regulatory Obligations Reference
Version	2.2
Baseline Date	2026-05-05
Classification	Internal — Restricted
Owner	Enterprise Architecture & Regulatory Engineering
Reviewers	CISO Office, Privacy Office, Model Risk Management, Third-Party Risk Management, Records & eDiscovery
Distribution	Internal use only. Not for distribution outside the firm without prior written approval from the document owner.
Review Cadence	Semi-annual; ad hoc on material change to the Microsoft Online Services Terms, the Microsoft Products and Services Data Protection Addendum (DPA), or the firm's Conditional Access posture.
Source Authorities	Microsoft Online Services Terms; Microsoft Products and Services DPA; Microsoft Trust Center; Microsoft Learn (M365 Copilot architecture, Privacy, Data Residency); 23 NYCRR Part 500; 17 CFR 240.17a-4; FINRA Rule 4511; Regulation (EU) 2022/2554 (DORA); Regulation (EU) 2016/679 (GDPR); Regulation (EU) 2024/1689 (EU AI Act); SR 11-7; OCC Bulletin 2011-12; SOX §404; MAS TRM Guidelines; HKMA SA-2; FCA SYSC 8; MiFID II Art. 16(7).

Revision History

Version	Date	Author	Summary
2.2	2026-05-05	Enterprise Architecture / Audit Defensibility	Audit-defensibility softeners across §2.2, §2.5, §2.7, §2.8 (CoWork), §2.10.1: (a) Substrate persistence qualified with “for Copilot experiences that support persistence, inheriting mailbox retention where applicable” — not all surfaces persist identically. (b) Anthropic non-retention reframed at first assertion as Microsoft contractual posture under the DPA + Product Terms (Enterprise Data Protection), not an architectural guarantee independently verifiable by the firm. (c) TLS / ExpressRoute language qualified with “where applicable / depending on tenant network design” — ExpressRoute terminates at the Microsoft edge and does not extend end-to-end to model providers. (d) Copilot Memories (IPM.Contact items) annotated with an emerging-behavior footnote — Microsoft documentation is still evolving. (e) CoWork “full request and response payloads” scoped to regulatory recordkeeping and supervisory review, otherwise minimized per firm policy — addresses GDPR Art. 5(1)(c) data minimization and insider-risk over-retention.

Version	Date	Author	Summary
2.1	2026-05-05	Enterprise Architecture	Labeling refinements for Legal and Audit clarity. Orchestrator now explicitly marked “Microsoft-managed routing — policy and safety enforced” throughout (reinforces that model choice does not bypass governance and that Anthropic is not directly user-addressable). “Content Safety / Prompt Filtering” relabeled as “Pre- and post-inference safety controls” to reflect the actual two-stage pipeline (safety before model call; safety after response, before user). AWS/GCP infrastructure naming intentionally omitted — regulators evaluate jurisdiction and contractual control, not hyper-scaler branding.
2.0	2026-05-05	Enterprise Architecture	Fact-check pass against Microsoft Learn and Trust Center sources. xAI reclassified from M365 Copilot sub-processor to independent processor (Copilot Studio integration only, US tenants); Microsoft DPA flow-down does not apply. Replaced “Zero Data Retention (ZDR)” with Microsoft’s actual term “Enterprise Data Protection (EDP)” — covered by the DPA and Product Terms. Removed unconfirmed claim that Anthropic operates on AWS/GCP — Microsoft has not publicly disclosed the underlying cloud platform. JWT lifetime corrected to 60–90 minutes default.
1.0	2026-05-05	Enterprise Architecture	Initial baselined release covering M365 E5 base architecture, the Copilot (Anthropic) overlay, and the Copilot CoWork orchestration layer. Replaces all prior working drafts.

Table of Contents

Executive Summary

This reference document explains, in two parts, the regulatory obligations carried by Microsoft 365 Enterprise E5 as deployed at the firm, and the incremental obligations introduced when Microsoft 365 Copilot (including the Anthropic model integration) and the firm's internal Copilot CoWork extension are layered on top. It is written for a Global Systemically Important Financial Institution (G-SIFI) audience and is intended to brief enterprise architects, identity leads, and compliance and risk officers on the controls already in place, the contractual protections provided by Microsoft, and the residual obligations that the firm itself must operate, evidence, and audit.

Part 1 establishes the base M365 E5 service architecture inside the Microsoft 365 service trust boundary, the identity, data, compliance, and telemetry planes, the data-residency posture, and the regulations that already attach to that base footprint. Part 2 overlays the Copilot architecture, identifies precisely where the Copilot data flow differs from native M365 data flow, describes the role of Anthropic as a Microsoft sub-processor, and explains where the Copilot CoWork orchestration layer participates. Each part concludes with a regulatory mapping table that ties services to specific obligations.

The central message of this document is straightforward. The Microsoft Products and Services Data Protection Addendum and the Microsoft Online Services Terms shift legal liability for many privacy and security obligations to Microsoft, but they do not, on their own, satisfy regulators. A G-SIFI must demonstrate active operational control: model risk governance, third-party risk monitoring, technical enforcement of data residency, immutable retention pipelines for prompt-and-response transcripts, access governance, and an independent audit trail. The compliance gap analysis in Part 3 sets out, in concrete terms, what the firm must continue to do regardless of contract.

Reading guidance

Part 1 is the platform reference; read it first if you have not deployed M365 E5.

Part 2 is the Copilot delta; read it first if you understand E5 and need only the incremental change.

Part 3 is the compliance gap analysis; read it before approving rollout.

Part 4 is a controls catalog mapped to regulations; use it for audit evidence and runbook construction.

Part 1 — Base M365 E5 Architecture and Existing Regulatory Obligations

1.1 Strategic Trust Boundary

The Microsoft 365 service trust boundary is the contractual and operational perimeter inside which tenant data is stored and processed under the Microsoft Online Services Terms and the Microsoft Products and Services DPA. Customer Data, as that term is defined in the DPA, is stored at rest in geographically defined regions, encrypted with Microsoft-managed keys (or customer-managed keys where the customer enables Customer Key), and processed only by Microsoft personnel and authorized sub-processors disclosed in the DPA's sub-processor list.

The Microsoft 365 service trust boundary is not the same boundary as Azure. Azure commercial cloud is a separate service family with its own DPA scope. Some M365 services rely on Azure-hosted components — most notably Azure OpenAI, which is the primary inference engine for M365 Copilot — but that reliance is governed by Microsoft's internal sub-processor flow-down, not by any direct customer contract with Azure. From the firm's perspective, the relevant commitments live in the M365 DPA and Online Services Terms, not in the Azure terms.

Boundary primitives

- **Control plane:** Identity, policy, and routing decisions (Microsoft Entra ID, Conditional Access, Purview policy engine, M365 Admin Center, tenant configuration).
- **Data plane:** User-generated content and metadata (Exchange Online mailboxes, SharePoint Online sites, OneDrive for Business, Teams chat and channel substrate, Loop, Stream).
- **Telemetry plane:** Service-generated activity records (Microsoft Purview Audit, Defender for Cloud Apps, Defender XDR advanced hunting, Entra sign-in and audit logs, Service Health).

Sub-processor disclosure model

Microsoft maintains a public sub-processor list under the DPA. New sub-processors are notified to customers in advance, and the customer has a contractual right to object. For M365 Copilot specifically, Microsoft has confirmed Anthropic as a sub-processor that may receive prompts, grounding context, and responses for inference purposes. xAI (Grok) is reachable from the Microsoft cloud only through the Copilot Studio integration as an independent processor — it is NOT a sub-processor under the M365 Copilot DPA, and Microsoft DPA flow-down does not apply. Use of xAI is therefore a firm-direct vendor relationship requiring separate contracts and TPRM review. Other model providers visible in the broader Azure AI Foundry catalog are not, today, confirmed M365 Copilot sub-processors with active DPA flow-down for M365 customer data; they are accessed under separate Azure terms.

Where data lives, by service

Service	Storage Location	Primary Compliance Stack
Exchange Online	Tenant-region mailbox database; with EUDB, EU/EFTA tenants pin to EU regions.	Purview Audit, Sensitivity Labels (encryption + IRM), DLP, Litigation Hold, Records Management, Communication Compliance.
SharePoint Online	Tenant-region content database; geo-targeted libraries via Multi-Geo.	Purview DLP, Sensitivity Labels, IRM, Records Management, eDiscovery (Premium).
OneDrive for Business	Per-user mysite, tenant-region by default.	Same as SharePoint; per-user retention; Insider Risk Management.
Microsoft Teams	Chat content in Substrate (Exchange-backed); files in SharePoint/OneDrive; meeting recordings in OneDrive/SharePoint Stream.	Communication Compliance, eDiscovery, Sensitivity Labels for chat, channel, and meetings.
Microsoft Loop	Tenant-region SharePoint Embedded container.	Sensitivity Labels, DLP, eDiscovery.
Substrate (hidden mailbox)	Per-user hidden Exchange mailbox folder used for system-generated content including Copilot interactions.	Tenant retention, Litigation Hold, eDiscovery, Communication Compliance, Records Management.

1.2 Identity Plane — Microsoft Entra ID and Conditional Access

Every M365 interaction begins with an authentication event. Microsoft Entra ID (formerly Azure AD) is the identity provider, the policy decision point for Conditional Access, and the issuer of access and refresh tokens. For E5 tenants, Entra ID P2 is included, which unlocks risk-based Conditional Access, Privileged Identity Management (PIM), and Identity Protection.

Token model

- **Access tokens:** JWT, default lifetime 60–90 minutes (configurable via Conditional Access token-lifetime policies); audience and scope are validated at every resource hop.

- **Refresh tokens:** Long-lived (default 90 days), bound to the device's Primary Refresh Token (PRT) when issued through the Web Account Manager (WAM) or MSAL; revocable via Entra session revocation.
- **On-Behalf-Of (OBO) flow:** A middle-tier API exchanges the user's token for a downstream-scoped token (for example, a Microsoft Graph token scoped to Mail.Read). OBO never elevates privileges; the downstream token always carries the original user's identity.

Conditional Access posture for a G-SIFI

1. Phishing-resistant MFA required for all users, evaluated at every sign-in (NYDFS §500.12).
2. Compliant device required for any access to sensitive workloads (Exchange Online, SharePoint, Teams, Copilot).
3. Sign-in risk and user risk evaluated continuously; high-risk sessions trigger step-up authentication or block.
4. Token Protection (where available) cryptographically binds tokens to the originating device.
5. Privileged role elevation gated through PIM with approval workflow and just-in-time activation.

> **Regulator's view**
>
> NYDFS §500.7 (Access Privileges and Management) requires that access privileges be granted based on the principle of least privilege and reviewed at least annually; for Class A companies, §500.7(c) additionally requires monitoring of privileged access activity and a privileged access management solution. Entra PIM, with logging into Defender XDR and the Unified Audit Log, is the primary technical control the firm uses to evidence compliance.

1.3 Data Plane

Customer Data in M365 E5 lives in four primary stores. Exchange Online holds mailbox content; SharePoint Online holds site and team content; OneDrive for Business holds personal files; Microsoft Teams stores chat content in a hidden Exchange mailbox (the Substrate) and files in SharePoint or OneDrive. Service-generated content from Copilot is also written into the Substrate, in a hidden folder structure that is discoverable through Purview eDiscovery.

The Substrate is, in regulatory terms, the most consequential storage location in M365. It is the eDiscovery anchor: anything written into the Substrate inherits the user's mailbox retention policies, can be placed on Litigation Hold, can be searched through eDiscovery (Premium), and can be exported for production. For a firm with SEC Rule 17a-4 or FINRA Rule 4511 obligations, Substrate retention is the primary mechanism by which Copilot interactions become discoverable records.

1.4 Compliance Plane — Microsoft Purview

Microsoft Purview is the umbrella name for the M365 compliance stack. The components most relevant to a G-SIFI are listed below; all are included in or licensable under E5.

Purview Component	Regulatory Function
Sensitivity Labels	Persistent classification and encryption of documents and emails. Labels survive download. Used to enforce confidentiality classification and Information Rights Management (IRM).
Data Loss Prevention (DLP)	Policy engine evaluating content against Sensitive Information Types (SITs) and trainable classifiers. Enforces blocking, justification, or override across Exchange, SharePoint, OneDrive, Teams, endpoints, and (with policy update) Copilot prompts and responses.
Records Management	Immutable record declarations with retention labels. The firm's primary native mechanism for SEC 17a-4(f) and FINRA 4511 immutable retention without an external WORM store.
eDiscovery (Premium)	Custodian-based legal hold, in-place search, review, and export. Discovers content across Exchange (including Substrate), SharePoint, OneDrive, Teams, and Copilot interactions.
Communication Compliance	Supervisory review of communications for code-of-conduct, market-abuse, and harassment patterns. Mandatory in scope for FINRA Rule 3110 supervisory review and similar regimes.
Insider Risk Management	Behavioral analytics over user activity (download spikes, unusual sharing, departing-user signals). Feeds Defender XDR.
Audit (Standard / Premium)	Unified Audit Log of tenant activity. Standard tier default retention 180 days; Premium tier default retention 1 year, extendable to 10 years with the add-on.
Data Lifecycle Management	Tenant-wide and label-scoped retention policies. Operates over all four data stores plus Substrate.
Information Barriers	Logical segmentation preventing communication or sharing between user groups; used for Chinese-wall enforcement in research/banking.
Data Security Posture Management (DSPM)	Discovery of sensitive-data exposure across the tenant; surfaces over-shared sites, broken inheritance, and risky permissions.

1.5 Telemetry Plane

Three independent telemetry streams cover the M365 estate. They are co-relatable but not yet fully unified, and the firm's SIEM integration is responsible for joining them on a common time, principal, and source-IP basis.

Stream	Source	Primary Records
Microsoft Purview Audit (UAL)	Exchange, SharePoint, OneDrive, Teams, Copilot, Entra-issued workloads.	Operation, UserId, ClientIP, Workload, ObjectId, AppId; high-fidelity activity records.
Microsoft Defender XDR	Defender for Identity, Defender for Cloud Apps, Defender for Office 365, Defender for Endpoint.	CloudAppEvents, IdentityLogonEvents, EmailEvents, DeviceEvents — KQL-queryable, joinable on AccountObjectId.
Microsoft Entra logs	Entra ID sign-in, audit, and provisioning events.	SignInLogs, AuditLogs, ProvisioningLogs; capture token issuance, Conditional Access evaluation, MFA outcomes, consent grants.

All three streams are exported into the firm's Microsoft Sentinel instance and forwarded to the enterprise SIEM. Correlation today relies on UPN, IP, and time-window pivots; native CorrelationId propagation continues to mature across workloads but is not yet end-to-end.

Audit retention — important nuance!

Microsoft Purview Audit (Standard) retains records for 180 days by default.

Microsoft Purview Audit (Premium) retains records for 1 year by default, with optional 10-year retention available as a paid add-on.

For a G-SIFI subject to SEC 17a-4 communications retention or DORA Article 12 record-keeping, audit retention alone is not the regulated record. The regulated record is the prompt or response itself, retained in the Substrate or in a parallel WORM store; UAL is the activity index, not the artifact.

1.6 Data Residency — EU Data Boundary and Multi-Geo

Microsoft operates the EU Data Boundary (EUDB) for European Union and European Free Trade Association (EFTA) customers. Under EUDB, Customer Data and pseudonymized personal data for the core M365 services are stored and processed within the EU/EFTA boundary. Service-generated diagnostic and security data has historically been an exception; Microsoft completed the inclusion of operations data and professional services data in stages through 2024–2025.

Multi-Geo, included in E5, allows the firm to pin specific users' Exchange, SharePoint, OneDrive, and Teams data to defined satellite regions outside the tenant's home region. This is the primary technical control for satisfying GDPR Chapter V transfer constraints, German BDSG residency expectations, and similar regional rules across the firm's footprint.

> **Regulator's view — residency**
>
> Contracts and Standard Contractual Clauses are not, on their own, technical controls. GDPR Chapter V, the EDPB's Schrems II guidance, and DORA Article 28(7) all expect demonstrable technical enforcement of where data is stored and processed. The firm's residency posture must be evidenced by Multi-Geo configuration, EUDB enablement, the M365 admin region map, and (where applicable) Customer Lockbox approvals — not just by the DPA Annex.

1.7 Regulatory Obligations Carried by Base M365 E5

The table below maps the regulations most material to a G-SIFI to the M365 E5 services and controls that bear on them. It is not exhaustive, and obligations vary by jurisdiction, business line, and entity registration; it is intended as the architecture-to-regulation cross-reference, not as legal advice.

Regulation	Obligation Type	M365 E5 Controls Engaged
GDPR (EU 2016/679)	Lawful basis, data subject rights, transfer controls, breach notification (72h).	EUDB, Multi-Geo, Sensitivity Labels, DLP, eDiscovery (DSAR fulfilment), Audit Premium, Defender XDR (breach detection), Customer Lockbox.
DORA (EU 2022/2554)	ICT risk management, incident reporting (4h/72h/1mo), critical ICT third-party register (Art. 28-30), threat-led penetration testing.	Defender XDR, Sentinel, Service Trust Portal evidence, Microsoft as critical ICT third party, Customer Lockbox audit.
EU AI Act (2024/1689)	Risk classification of AI systems; transparency, logging, human oversight, post-market monitoring for high-risk uses.	(Engaged primarily in Part 2 — Copilot overlay.) Documented limitations, audit logs of AI interactions, DPIA artefacts.
NYDFS 23 NYCRR Part 500	Cybersecurity program, governance, MFA, third-party policy, asset/data retention, incident response, 72h notification.	Entra PIM, Conditional Access, MFA, Sensitivity Labels + DLP, Defender XDR, Audit Premium, Records Management.

Regulation	Obligation Type	M365 E5 Controls Engaged
SEC Rule 17a-4 / 18a-6	Books-and-records preservation: 6 years (first 2 readily accessible) for most records; 3 years for communications under 17a-4(b)(4); WORM/audit-trail electronic format.	Records Management, Litigation Hold, eDiscovery, immutable retention labels, third-party WORM archives where required.
FINRA Rule 4511	General recordkeeping defaulting to 6 years; ties to 17a-4 storage formats.	Same Purview Records and eDiscovery stack as 17a-4.
FINRA Rule 3110	Supervisory review of communications.	Communication Compliance, Insider Risk, supervisory review workflows.
SOX §404	Internal Controls over Financial Reporting (ICFR) — IT general controls.	Entra PIM, Conditional Access, segregation-of-duty enforcement, Audit Premium evidence, change-management logging.
MiFID II Art. 16(7)	Recording of telephone and electronic communications; retention 5 years (extendable to 7).	Teams call recording, Communication Compliance, Records Management.
MAS TRM, HKMA SA-2, FCA SYSC 8	Outsourcing risk management, operational resilience, supervisory access.	Microsoft Service Trust Portal, attestations (SOC 2 Type II, ISO 27001/27017/27018, PCI), Customer Lockbox, exit-strategy clauses.
GLBA Safeguards Rule	Administrative, technical, and physical safeguards for non-public personal information.	Sensitivity Labels (NPI SITs), DLP, encryption at rest and in transit, Defender XDR.
CCPA/CPRA, NYDFS Cybersecurity, state breach laws	Notification timelines, NPI handling.	DLP, eDiscovery, Audit Premium, Defender XDR.

1.8 Where Microsoft's Contract Ends

The DPA and Microsoft Online Services Terms commit Microsoft to a defined set of confidentiality, security, and processing obligations as a processor (and, in narrow cases, as an

independent controller). They do not, and cannot, satisfy the firm's obligations as the controller. The firm remains responsible for governance, lawful basis, data subject rights handling, retention configuration, model risk management decisions, supervisory review, and the design and operation of compensating controls. Part 3 of this document develops the gap analysis in detail.

Part 2 — Microsoft 365 Copilot Overlay (Anthropic) and Copilot CoWork

2.1 What Copilot Adds, in One Picture

Microsoft 365 Copilot is a generative-AI experience that operates inside the M365 service trust boundary and uses the user's existing Microsoft Graph permissions to retrieve grounding context. It is licensed as a per-user add-on to E5 (or as part of the Microsoft 365 Copilot SKU). Architecturally, it introduces a new processing component — the Copilot orchestrator — between the user and the existing M365 data plane, and it introduces a new sub-processor (Anthropic) that may receive prompts and grounding context for inference under the M365 Copilot DPA flow-down. xAI (Grok) is NOT a sub-processor on this path; it is reachable only through the Copilot Studio integration as an independent processor, governed by the firm's own vendor relationship with xAI rather than the Microsoft DPA.

Copilot CoWork is the firm's internal extension layer. It is not a Microsoft service. CoWork sits between the user's M365 client and the Copilot orchestrator (or, depending on deployment posture, alongside it as a parallel agentic service surfaced through the M365 Gateway). Its purpose is to broker On-Behalf-Of tokens, enforce additional firm-specific policy, and run agentic skills that the firm has authored. From a regulatory perspective, CoWork does not change the M365 service boundary, but it does add a firm-operated processing layer that is the firm's responsibility to govern, log, and validate.

2.2 The Copilot Orchestrator

The Copilot orchestrator is the M365 service component that receives a user prompt, performs prompt safety filtering, retrieves grounding data from Microsoft Graph using the user's permissions (semantic index, recent files, calendar, mail, chats), constructs the inference payload, routes the inference to a model, post-processes the response (DLP evaluation, label propagation, content safety), and returns the grounded response with inherited classification metadata.

Internal stages

- **Pre-processing:** Prompt-safety filtering, jailbreak detection, prompt-injection heuristics.
- **Grounding:** Microsoft Graph queries scoped to the user's identity; the user never sees data they could not already access. Sensitivity labels on grounding documents are read and inherited.
- **Model routing:** By default, Azure OpenAI; with admin opt-in and geographic eligibility, Anthropic. Routing decisions are evaluated against tenant admin controls and the user’s region. xAI is not in this routing set — Grok is reachable only via Copilot Studio as an independent processor, not via the M365 Copilot inference path.

- **Post-processing:** DLP evaluation of the response, content-safety scoring, citation construction, sensitivity-label propagation to any artefacts produced.
- **Persistence:** For Copilot experiences that support persistence, prompt and response are written to a hidden folder in the user's Substrate mailbox, inheriting the user's mailbox retention where applicable. Where persisted, they are searchable through eDiscovery and can be placed on Litigation Hold. Not all Copilot surfaces persist identically — surface-by-surface behavior should be confirmed against current Microsoft documentation.

2.3 Model Provider Landscape

Model Provider	Sub-Processor Status	Geographic Eligibility	Default Posture
Azure OpenAI	Microsoft (intra-company); operated under Azure, flowed down through the M365 DPA.	Global, including EUDB.	Enabled by default; primary inference path.
Anthropic (Claude)	Confirmed sub-processor under the Microsoft DPA for M365 Copilot.	Disabled by default in EU, EFTA, and UK; unavailable in government and sovereign clouds; enabled by default in US, with admin opt-out.	Admin opt-in/opt-out per tenant; routing controllable through M365 Admin Center.
xAI (Grok)	Independent processor — NOT a sub-processor under the M365 Copilot DPA. Reachable from the Microsoft cloud only via the Copilot Studio integration; firm vendor relationship governs.	US tenants only.	Admin opt-in; not available outside US.
Other Foundry models (DeepSeek, Mistral, Llama, Qwen, etc.)	Available in the broader Azure AI Foundry catalog under Azure terms; not, today, confirmed M365 Copilot sub-processors with active DPA flow-down for M365 customer data.	Subject to per-model Azure region availability.	Out of scope for M365 Copilot routing; if used, accessed through separate Azure tenancy and contracts.

Model Provider	Sub-Processor Status	Geographic Eligibility	Default Posture

Regulator's view — sub-processors

DORA Article 28(2) requires the firm to maintain a register of all contractual arrangements on the use of ICT services provided by ICT third-party service providers. Anthropic, when enabled, is a sub-processor of Microsoft, not a direct contracting party of the firm; it must nonetheless appear in the register as a sub-processor with the chain of accountability explicit.

NYDFS §500.11 requires a written third-party service provider security policy and due diligence; the firm's policy must address sub-processor changes and contemplate Microsoft's right to add or change sub-processors with notice.

2.4 Anthropic Integration — Specific Mechanics

When Anthropic is enabled, Copilot may route inference to Claude through a Microsoft-managed connection to Anthropic. The connection is governed by the M365 DPA and Microsoft's contract with Anthropic, not by any direct contract between the firm and Anthropic. The data sent comprises the user's prompt, grounding context retrieved from the user's authorized Graph corpus, and any system instructions; the data returned is the model's response. Microsoft's contractual posture is that Anthropic acts as a sub-processor of Microsoft, governed by Microsoft's Enterprise Data Protection (EDP) commitments under the DPA and Product Terms: Anthropic does not train its models on Customer Data and does not retain it beyond inference.

Enterprise Data Protection is contractual, not architectural.

Enterprise Data Protection (EDP) is Microsoft's term for the contractual no-training and non-retention commitments applied to Customer Data under the DPA and Product Terms. It is evidenced by Microsoft's vendor controls and the sub-processor's audit programme. It is not enforced by hardware or by a cryptographic mechanism on the customer's premises. Regulators distinguish between contractual residency (covered by the DPA Annex) and technical residency (evidenced by routing controls, EUDB enablement, and admin-policy enforcement). The firm should treat EDP as a vendor-management control, monitor it through TPRM, and not assume it removes the obligation to perform a DPIA.

Geographic restriction enforcement

For EU/EFTA/UK tenants, Anthropic is disabled by default. Re-enabling requires an explicit administrative action, and the firm's Conditional Access and admin-RBAC posture must restrict who can perform that action. A tenant-level setting, monitored through Audit Premium, records any change. The firm's recommended posture is: keep Anthropic disabled in EU tenants pending DPIA completion, and pin the M365 Admin Center setting through Microsoft Defender for Cloud Apps Conditional Access App Control where available.

2.5 Copilot CoWork — The Firm's Extension Layer

Copilot CoWork is the firm's internal product layer that surfaces curated agentic skills to users, brokers downstream API calls under the user's identity, and enforces firm-specific policy beyond what Microsoft's native Copilot policy supports. CoWork is a firm-operated workload; it inherits the firm's existing identity, telemetry, and compliance controls and is governed under the firm's standard Software Development Life Cycle, change management, and operational risk processes.

CoWork's relationship with M365 and Copilot

- **Identity:** CoWork is a registered Entra ID application. It accepts the user's M365 token and performs OBO exchange against Microsoft Graph and other downstream APIs, never elevating privileges and always operating within the user's effective Graph permissions.
- **Trust boundary:** CoWork is hosted in the firm's Azure tenancy, in regions selected to align with the firm's residency posture. It is inside the firm's perimeter; from Microsoft's perspective, CoWork is a customer-operated application, not part of the M365 service.
- **Data flow:** User prompt → CoWork orchestrator (firm-operated) → Microsoft Graph (via OBO) → optional model invocation (Azure OpenAI in firm tenancy, or M365 Copilot via the M365 Gateway pattern) → CoWork policy/DLP layer → response to user.
- **Logging:** Every CoWork prompt, every downstream call, every model invocation, and every response is written to the firm's SIEM — full request and response payloads where required for regulatory recordkeeping and supervisory review, otherwise minimized per firm policy. Persisted on the firm's WORM store under the 17a-4 / FINRA 4511 retention schedule, with the minimization tier reconciling those obligations against GDPR Art. 5(1)(c) and insider-risk over-retention concerns.

> **CoWork is the firm's responsibility, end-to-end.**
>
> Microsoft's DPA does not cover content that flows through CoWork; CoWork is the firm's processor.
>
> If CoWork invokes a third-party model (for example, Anthropic directly under an Anthropic-customer contract, not via Microsoft), that vendor relationship is the firm's, governed by the firm's TPRM and DPIA framework.
>
> Where CoWork surfaces M365 Copilot through the M365 Gateway pattern, the data path inside Copilot stays under the M365 DPA; the data path inside CoWork is still the firm's responsibility. The two boundaries meet at the Gateway.

2.6 OBO Token Mechanics in the Copilot/CoWork Path

On-Behalf-Of is the OAuth 2.0 token exchange flow that allows a middle-tier service to call a downstream API as the originating user. Both native Copilot and CoWork rely on OBO when

calling Microsoft Graph for grounding context. The mechanics are well documented; the firm's responsibility is to enforce the policy posture below.

OBO Element	Firm Policy Requirement
Audience	All middle-tier APIs validate audience strictly. Tokens minted for one resource must never be replayed against another.
Scopes	Downstream tokens must be scoped to the minimum required (Mail.Read, Sites.Read.All, Files.Read.All, etc.). Application permissions are forbidden in the Copilot/CoWork user-context path.
Claims	oid, tid, ipaddr, device_id, and amr (authentication-method reference) validated at every hop; mismatched claims trigger token rejection.
Lifetime	Default 60–90 minutes; shorter lifetimes available via Conditional Access token-lifetime policies for high-sensitivity workloads.
Refresh	Refresh tokens device-bound through PRT/WAM; loss-of-device triggers revocation through Entra session revocation.
Risk re-evaluation	Mid-session risk-score change re-evaluates Conditional Access; high-risk state blocks the OBO exchange and requires step-up authentication or session termination.

2.7 Storage and Discoverability of Copilot Interactions

M365 Copilot prompts and responses are persisted to the user's Substrate (hidden mailbox folder). They inherit the user's Exchange mailbox retention policy and Litigation Hold status, are searchable through Purview eDiscovery (Standard or Premium), and are exportable for production. This makes Copilot interactions discoverable records in the same way that emails are; for a G-SIFI subject to SEC 17a-4, FINRA 4511, MiFID II Art. 16(7), or DORA Article 12 record-keeping, this is the principal mechanism by which Copilot conversations enter the regulated record.

CoWork interactions are stored separately. The firm's CoWork pipeline writes the prompt, the grounding context, the model invocation metadata, and the response to a structured log in the firm's SIEM, and forwards a copy to the WORM archive. Retention is set to the firm's communications retention schedule (typically 7 years to provide a safety margin over the 6-year statutory 17a-4 minimum, plus jurisdictional adders where applicable).

Records retention — practical

SEC Rule 17a-4(a) requires 6 years for most books and records, the first 2 readily accessible.
SEC Rule 17a-4(b)(4) requires 3 years for communications received and copies of

communications sent (with first 2 readily accessible).

FINRA Rule 4511 default is 6 years; many firms adopt 7 years internally as a safety margin.

MiFID II Art. 16(7) requires 5 years (extendable to 7 by competent authorities).

DORA Article 12 requires record-keeping of ICT-related incidents for at least 5 years.

Set the firm's Copilot and CoWork retention to the longest applicable horizon across the jurisdictions in scope, then apply jurisdictional shorter holds only where a regulator specifically requires deletion.

2.8 Telemetry Surface for Copilot and CoWork

Stream	Copilot Coverage	CoWork Coverage
Purview Audit (UAL)	Copilot Interaction operations: Workload=Copilot, with prompt timestamp, app context, mailbox/file references, and (where enabled) prompt/response identifiers.	Not native; CoWork emits structured equivalents into the firm's SIEM and references UAL CorrelationId where available.
Defender XDR — CloudAppEvents	App activity for the M365 Copilot app, file access events, DLP hit events.	App activity for the registered CoWork application; file access through OBO is captured against the user's account.
Entra Sign-In Logs	Token issuance for the Copilot first-party app; CA evaluation results; MFA outcomes.	Token issuance for the CoWork app registration; OBO exchange events; consent grants.
Substrate (mailbox-resident records)	Prompt and response artefacts retained as the user's records.	Not used; CoWork uses the firm's WORM store.

2.9 Incremental Regulatory Obligations Introduced by Copilot

The base M365 E5 obligations identified in §1.7 continue to apply unchanged. The table below identifies the additional obligations that attach when Copilot (with Anthropic) and CoWork are introduced, and the controls engaged to meet them.

Regulation / Domain	Incremental Obligation	Controls Engaged
EU AI Act (2024/1689)	Provider/deployer obligations: GPAI transparency, Article 26 deployer duties (instructions, human oversight, monitoring), risk classification of any high-risk use case the firm builds on top.	DPIA per Copilot use case, model-card review, human-in-the-loop configuration, output logging in CoWork, EU AI Office register where required.
SR 11-7 / OCC 2011-12	Model risk identification and validation for use cases that influence financial reporting, trading, or compliance decisions.	AI Risk Committee, model inventory entry per use case, independent validation, drift monitoring, documented limitations.
NYDFS §500.11 (Third-Party Service Provider Security Policy)	Inclusion of Anthropic (and any other model sub-processor) in the third-party register; due diligence on Microsoft's vendor controls over Anthropic.	TPRM intake for sub-processor changes, attestation review (SOC 2 Type II, ISO 27001), Customer Lockbox monitoring.
NYDFS §500.13 (Asset Management & Data Retention)	Inclusion of Copilot interactions in the asset inventory and data retention policy.	Records Management labels covering Substrate Copilot folders; CoWork retention schedule.
NYDFS §500.16 / §500.17	Incident response plan covering Copilot-induced incidents; 72-hour notification to the Superintendent for qualifying incidents.	Defender XDR runbooks, Sentinel automation, legal/compliance notification workflow.
DORA Articles 17–23 (Incident Reporting)	Major ICT-related incident reporting: initial notification ≤4h after classification (≤24h after detection), intermediate ≤72h, final ≤1 month.	Sentinel-driven incident classification, reporting workflow into the National Competent Authority portal, ESA forwarding.

Regulation / Domain	Incremental Obligation	Controls Engaged
DORA Articles 28–30 (Critical ICT Third Parties)	Microsoft and (transitively) Anthropic appear in the firm's ICT third-party register; designation of Microsoft as a Critical ICT Third-Party Provider (CTPP) by the ESAs triggers additional oversight.	Register maintenance, exit strategy documentation, threat-led penetration testing where required, concentration risk assessment.
GDPR Chapter V & EDPB Schrems II	Transfer impact assessment for Anthropic where enabled outside EU; technical residency enforcement.	DPIA, EUDB enablement, Multi-Geo, admin policy disabling Anthropic in EU/EFTA/UK by default.
SEC 17a-4 / FINRA 4511	Copilot and CoWork prompt/response records subject to communications recordkeeping.	Substrate retention via Records Management; CoWork WORM pipeline; Litigation Hold extension to Copilot interactions.
FINRA Rule 3110	Supervisory review of Copilot-generated content where it constitutes a communication.	Communication Compliance policies covering Copilot content; supervisory queues.
MiFID II Art. 16(7)	Where Copilot is used to generate communications relating to client orders/transactions, those communications fall within recording scope.	Records Management retention, Communication Compliance, eDiscovery.
MAS TRM, HKMA SA-2, FCA SYSC 8 (Outsourcing)	Material outsourcing assessment for Copilot; supervisor notification where applicable; right of audit through the Service Trust Portal.	Outsourcing register, regulator notification workflow, SOC 2 / ISO attestations.

2.10 End-to-End Data Flow Reference

The sub-sections below sequence the canonical data flows that a single Copilot interaction traverses, from the user's keystroke to the persisted record. They are written as architect-readable step tables rather than diagrams; each step identifies the trust boundary in force, the action performed, and the firm controls that must be evidenced at that hop. The flows are

observable in Microsoft Purview Audit, Defender XDR, and Entra sign-in logs, and should reconcile end-to-end on a common CorrelationId where the workload emits one.

2.10.1 Eleven-Step Copilot Prompt Lifecycle

The lifecycle below traces a single Copilot Chat prompt from a managed device to a persisted Substrate record. It crosses three trust boundaries: the user/device perimeter, the M365 service trust boundary, and (where Anthropic is enabled) the model sub-processor perimeter.

#	Boundary / Hop	Action	Controls Engaged
1	User / Managed Device	User submits prompt from a Copilot-enabled M365 surface (Word, Teams, Outlook, Edge, Copilot Chat).	Conditional Access compliant-device check; device-bound Primary Refresh Token (PRT); local DLP endpoint policies.
2	Device → WAM/MSAL	Identity broker presents PRT and requests an access token for the Copilot resource.	WAM token cache; TPM-backed device key; phishing-resistant MFA evaluation if required.
3	WAM/MSAL → Entra ID	Entra evaluates Conditional Access — sign-in risk, user risk, named locations, MFA strength, app-protection policy.	Entra ID P2 risk evaluation; Conditional Access; PIM where role activation is involved.
4	Entra ID → Client	JWT access token issued, audience-bound to the Copilot first-party app, scoped to the user's effective Graph permissions.	Token lifetime policy; Token Protection (binding) where enabled; sign-in log written to Entra.
5	Client → Copilot Orchestrator	Prompt and token submitted to the orchestrator over TLS within the M365 service trust boundary.	TLS 1.2+; mutual TLS for service-to-service; audience and scope validation; UAL Copilot Interaction event begins.
6	Orchestrator → Pre-processing	Prompt-safety filtering, jailbreak heuristics, prompt-injection detection.	Microsoft content-safety service; prompt log written to telemetry stream; UAL event enriched.
7	Orchestrator → Microsoft Graph (OBO)	On-Behalf-Of exchange yields downstream-scoped tokens; Graph queries retrieve grounding context (mail, files, chats, calendar) under the user's permissions.	OBO scope minimizations; sensitivity labels read on every grounding artefact; Graph access logged in CloudAppEvents.

#	Boundary / Hop	Action	Controls Engaged
8	Orchestrator → Model Routing	Microsoft-managed routing — policy and safety enforced. The orchestrator (not the user) selects the model based on tenant admin policy and user region: Azure OpenAI by default; Anthropic where opted in and geographically eligible. Anthropic is not directly user-addressable. xAI is not in this routing set — Grok is reached only via Copilot Studio as an independent processor, outside the M365 Copilot DPA flow-down.	Admin model-pinning policy; Defender for Cloud Apps Conditional Access App Control over admin center; Audit Premium records change events.
9	Model boundary crossing	Inference payload (prompt + grounding context + system instructions) sent to selected model. For Anthropic: traffic crosses into Anthropic's infrastructure under Microsoft DPA flow-down with Microsoft's contractual Enterprise Data Protection commitments.	DPA flow-down; Enterprise Data Protection contractual clause; sub-processor register entry; DPIA where required; vendor SOC 2 Type II review.
10	Model → Orchestrator (response)	Response returned; post-processing performs DLP evaluation, content-safety scoring, citation construction, and sensitivity-label propagation to any artefacts produced.	Purview DLP policy; Purview Information Protection label inheritance; Communication Compliance hooks.
11	Orchestrator → User + Substrate	Response delivered to the user; for Copilot experiences that support persistence, prompt and response are written to a hidden folder in the user's Substrate mailbox, inheriting mailbox retention where applicable.	Records Management label; Litigation Hold (where applied); eDiscovery indexing; UAL CopilotInteraction event closes with full identifiers.

Three boundaries, three accountabilities

Steps 1–4 sit in the user/device perimeter — the firm is fully accountable.

#	Boundary / Hop	Action	Controls Engaged

Steps 5–8 and 10–11 sit inside the M365 service trust boundary — Microsoft is the processor under the DPA.

Step 9 is the only step that crosses into a sub-processor's infrastructure (Anthropic). Microsoft remains the contracting processor; the sub-processor operates under DPA flow-down. The firm's TPRM register must reflect this transitive relationship explicitly. xAI is not on this path — it is reached only through Copilot Studio as an independent processor, governed by the firm's own vendor contract.

2.10.2 Authentication Token Chain

The token chain below is the OAuth 2.0 / OIDC flow that issues the access token used in step 4 of the prompt lifecycle. It is the firm's primary technical evidence of NYDFS §500.12 (MFA), SOX §404 ITGC (authentication and access), and ISO 27001 A.9 obligations.

#	Node	Function	Material Claims / Evidence
1	Device (TPM + PRT)	Windows broker holds the Primary Refresh Token, sealed against the TPM and bound to the device's hardware identity.	device_id claim; TPM attestation; Conditional Access compliant-device evaluation.
2	WAM / MSAL	Web Account Manager or MSAL libraries broker token requests, present the PRT, and cache resulting access/refresh tokens.	Refresh-token rotation; cache encryption at rest under DPAPI.
3	Entra ID + Conditional Access	Authoritative identity provider; evaluates user risk, sign-in risk, named locations, MFA strength, and app-protection policy.	Sign-in log; Conditional Access result; risk score.
4	JWT Issued	Access token returned: audience-bound, scope-restricted, lifetime 60–90 minutes by default.	aud, scp, oid, tid, ipaddr, amr, device_id claims; signed by Entra signing key.

#	Node	Function	Material Claims / Evidence
5	OBO Exchange	Middle-tier service (Copilot orchestrator, CoWork) exchanges the user's token for downstream-scoped tokens (Graph, third-party APIs) without elevating privileges.	OBO assertion; downstream audience and scope; user identity preserved end-to-end.
6	Active Session	Resource APIs validate the token at every hop; mid-session risk-score change re-evaluates Conditional Access; Token Protection binds tokens to the originating device where enabled.	Per-resource token validation; revocation via Entra session revocation; UAL records token-issuance and consent events.

2.10.3 Grounded Prompt Assembly

Grounding is the retrieval phase in which the orchestrator pulls context from the user's authorized Graph corpus and constructs the inference payload. It is the step at which the user's effective permissions become the upper bound of Copilot's data exposure; over-shared sites and broken inheritance directly translate into over-grounded prompts.

#	Source / Component	Action	Compliance Hook
1	Microsoft Graph — Mail (Exchange Online)	Retrieves recent mail items relevant to the prompt under Mail.Read scope on the user's behalf.	Sensitivity labels on messages are read; encrypted (IRM-protected) mail is decrypted only inside the M365 trust boundary.
2	Microsoft Graph — Files (SharePoint / OneDrive)	Retrieves file content under Files.Read.All / Sites.Read.All scope as effective on the user's identity.	Document sensitivity labels read; restricted-label content excluded from response surfacing if policy demands.
3	Microsoft Graph — Chats (Teams)	Retrieves recent chat and channel messages relevant to the prompt under Chat.Read / ChannelMessage.Read.All scope.	Communication Compliance retains visibility on Copilot-derived chat content.

#	Source / Component	Action	Compliance Hook
4	Microsoft Purview — Label and DLP Check	Each retrieved artefact is evaluated against the tenant's DLP policy; sensitivity labels are merged for the response.	DLP block / justification / override events written to UAL; label inheritance recorded for audit.
5	Orchestrator — Grounded Prompt Assembly	Filtered, label-aware grounding context is concatenated with system instructions and the user prompt to form the inference payload.	Token-budget enforcement; prompt-injection heuristics applied to grounding text.
6	Orchestrator — Token Count and Routing	Final payload size and routing decision logged; payload handed to the selected model under §2.10.4.	UAL CopilotInteraction event includes payload size and model route; DSPM surfaces over-grounding patterns.

2.10.4 Inference Path and Enterprise Data Protection Layered Analysis

The inference flow below begins where §2.10.3 ends. It is the most regulator-sensitive part of the lifecycle because it is the only segment that may cross into a sub-processor's infrastructure.

#	Stage	Action	Controls Engaged
1	Orchestrator → Pre-inference safety controls	Pre-inference stage of Microsoft's two-stage safety pipeline. Prompt and grounding context evaluated by Microsoft content-safety service before the model call; the post-inference stage screens the response before it returns to the user.	Microsoft safety-classifier outputs logged.
2	Orchestrator → Purview Pre-Hook	DLP and information-protection policies evaluated against the assembled payload.	DLP outcome written to UAL.
3	Orchestrator → Model Selector	Routing decision: Azure OpenAI (default) or Anthropic (where eligible). xAI is not on this path — Grok is a Copilot Studio integration, independent processor.	Admin policy; Defender for Cloud Apps CA App Control; geographic rule evaluation.

#	Stage	Action	Controls Engaged
4	Crossing the model boundary	Inference payload dispatched to selected model over a Microsoft-managed connection.	DPA flow-down; sub-processor register entry; transit encryption.
5	Model inference	Model returns response. Per Microsoft's contractual commitments under the DPA and Product Terms (Enterprise Data Protection), Anthropic does not train on Customer Data and does not retain it beyond inference. This is contractual posture, not an architectural guarantee independently verifiable by the firm.	Enterprise Data Protection contractual clause; vendor SOC 2 Type II; firm DPIA.
6	Orchestrator → Purview Post-Hook	Response evaluated against DLP, content-safety, and label-propagation policies.	Sensitivity-label propagation to artefacts produced by the response.
7	Orchestrator → User + Substrate	Response surfaced to user; for Copilot experiences that support persistence, prompt and response are written to the Substrate hidden folder, inheriting mailbox retention where applicable.	Records Management; Litigation Hold; eDiscovery indexing.

Enterprise Data Protection — layered confidence model

Enterprise Data Protection (EDP) is not a single control. It is a stack of four layers, each with a different enforcing party and different evidence quality. The firm should report EDP confidence at the layer level, not as a binary, and should treat architectural-layer enforcement as the lowest-confidence band given the limits of independent verification.

Layer	Mechanism	Enforced By	Evidence / Confidence
Contractual	Microsoft DPA and Online Services Terms; sub-processor flow-down clauses.	Legal agreement.	DPA audit clause; high confidence.

Layer	Mechanism	Enforced By	Evidence / Confidence
Operational	Sub-processor data-handling procedures and personnel controls.	Vendor SOC 2 Type II programme.	Annual audit reports; medium confidence.
Architectural	Stateless inference endpoint; no persistent storage of payload at the model provider.	Vendor infrastructure design.	Not independently verifiable from the firm's side; low confidence.
Monitoring	Unified Audit Log and eDiscovery inside M365 capture every Copilot interaction.	Microsoft Purview.	Admin-accessible; high confidence.

Reporting Enterprise Data Protection honestly to regulators

When the firm reports its Enterprise Data Protection posture to a regulator, it should distinguish the four layers. The contractual and monitoring layers are high-confidence and well-evidenced. The operational layer rests on third-party attestation. The architectural layer cannot be independently verified by the firm. Regulators familiar with model-provider architectures (notably the EDPB and the ESAs under DORA) expect this nuance — overclaiming “EDP enforced” without layering invites a finding of inadequate vendor oversight.

2.10.5 Logging and Retention Channels

Copilot interactions are visible across three independent logging channels. Each captures different content, has different default retention, and is accessed through a different administrative path. The firm's recordkeeping posture must explicitly identify which channel holds the regulated record under SEC 17a-4, FINRA 4511, MiFID II Art. 16(7), and DORA Article 12.

Channel	Content Captured	Default Retention	Access Method	17a-4 Posture
UAL (Microsoft Purview Audit)	Activity metadata: timestamp, user, app, mailbox/file references, prompt and response identifiers.	Standard 180 days; Premium 1 year (10 years with add-on).	Purview Audit search; SIEM export.	Activity index only — not the regulated record.

Channel	Content Captured	Default Retention	Access Method	17a-4 Posture
Substrate (Exchange hidden mailbox)	Full prompt and full response content, for Copilot experiences that support persistence, in the user's hidden mailbox folder.	User mailbox retention policy; Litigation Hold extends.	Purview eDiscovery (Standard or Premium).	The regulated record. Configure retention to the longest applicable horizon.
Copilot Memories (IPM.Contact items) †	User-curated memory items the model uses for personalization.	Indefinite by default; user-managed.	User UI; eDiscovery surfaces them as items.	User-generated content; subject to communications retention if used in regulated workflows.

† Copilot Memories (IPM.Contact items): based on current Microsoft documentation and observed behavior as of the baseline date. Microsoft's public documentation on Copilot memory storage, item class, and discoverability is still evolving — treat as an emerging behavior and reconfirm against current Microsoft Learn / Trust Center guidance before relying on it for a specific regulatory finding.

The dual-path design

UAL and the Substrate are deliberately separated. UAL is the activity index; the Substrate holds the artefact. A request from a regulator for "the records of all Copilot interactions involving custodian X for matter Y" is satisfied by a Substrate eDiscovery search, not by a UAL extract. Build the firm's eDiscovery runbook around this distinction.

Part 3 — Compliance Gap Analysis

3.1 What the DPA and MBSA Cover

The Microsoft Products and Services Data Protection Addendum and the Microsoft Master Business and Services Agreement (MBSA), together with the Microsoft Online Services Terms, commit Microsoft to a specific and material set of obligations as a processor: defined security controls, sub-processor disclosure, breach notification, audit support, support for data subject rights requests, GDPR Article 28 processor warranties, EU SCCs and UK IDTA where applicable, and Standard Contractual Clauses for international transfers. The DPA includes a Data Subject Rights addendum, a Sub-Processor list, and the EUDB commitments.

The MBSA shifts liability for Microsoft's own breaches and provides a contractual basis for indemnification within the MBSA's negotiated caps. The Service Trust Portal exposes the third-party attestations (SOC 1 Type II, SOC 2 Type II, ISO 27001/27017/27018, FedRAMP, HIPAA where applicable, PCI DSS where applicable, IRAP, C5, ENS) that allow the firm to evidence its supplier due-diligence position to its own regulators.

3.2 What the DPA and MBSA Do Not Cover

The DPA and MBSA cannot, by their nature, satisfy the firm's obligations as the controller. Regulators audit the firm's governance, not Microsoft's contracts. The five domains below set out where the firm must continue to operate compensating controls.

3.2.1 Model Risk Management and AI Governance

Authorities: SR 11-7 (Federal Reserve), OCC Bulletin 2011-12, EU AI Act (2024/1689), MAS Notice on AI in Financial Services (where in force), FCA AI guidance.

The DPA does not cover model selection, model validation, or output verification. The firm must establish an AI Risk Committee with a documented charter, maintain a model inventory entry per Copilot or CoWork use case, perform independent validation proportionate to materiality, monitor model drift, and document model limitations. For use cases that influence trading, compliance, or financial reporting, the firm must additionally enforce Human-in-the-Loop (HITL) review.

3.2.2 Third-Party Risk Management and Exit Strategy

Authorities: DORA Articles 28–30 (EU), NYDFS §500.11, MAS TRM Outsourcing, HKMA SA-2, FCA SYSC 8.

Microsoft's DPA does not perform the firm's third-party risk assessment for it. The firm must maintain a register of all ICT third-party arrangements, including Anthropic and other Microsoft sub-processors flowing down through the DPA. The register must capture criticality, concentration risk, exit strategy, and recovery time objectives. A Copilot Decommissioning

Runbook — documented, exercised at least annually, and approved at the appropriate operational risk committee — is part of the firm's audit evidence.

3.2.3 Technical Enforcement of Data Residency

Authorities: GDPR Chapter V, EDPB Schrems II guidance, DORA Article 28(7), MAS TRM 11.0.5 (data sovereignty), HKMA SA-2 Annex.

Standard Contractual Clauses are not technical controls. The firm must enforce residency through M365 Admin Center settings (EUDB, Multi-Geo), Conditional Access Country/Named Locations, model-routing pinning (Anthropic disabled in EU tenants pending DPIA), and Customer Lockbox approvals for any Microsoft engineer access. Evidence of enforcement — exported configuration, screenshots, change logs — is the firm's responsibility.

3.2.4 eDiscovery, Recordkeeping, and WORM Retention

Authorities: SEC Rule 17a-4 (17 CFR 240.17a-4), FINRA Rule 4511, MiFID II Art. 16(7), DORA Article 12.

Purview Audit (UAL) is not the regulated record. UAL is the activity index. The regulated record is the prompt or response itself. The firm must configure Records Management labels with retention and disposition over the Substrate hidden folders that store Copilot interactions, place all relevant custodians on Litigation Hold during active matters, and operate a parallel WORM pipeline for CoWork interactions and any logs that fall outside Substrate. A Sentinel-to-Logic-App-to-WORM pipeline is the typical pattern; native Records Management is the preferred path where it covers the artefact.

3.2.5 Access Governance and Least Privilege

Authorities: NYDFS §500.7, §500.12, SOX §404 ITGC, ISO 27001 A.9, COBIT DSS05.

Broad delegated scopes (Sites.Read.All, Files.Read.All) make Copilot's grounding surface as wide as the user's effective Graph permissions. The firm must enforce phishing-resistant MFA and compliant device through Conditional Access targeted at the Copilot and CoWork applications, monitor Graph permission exposure through DSPM, and run quarterly access reviews through Entra Access Reviews on any Privileged Identity Management roles that can change Copilot tenant policy.

3.3 The Final Verdict

Domain	What DPA/MBSA Provides	What the Firm Must Still Do
Liability	Legal indemnification within negotiated caps for Microsoft-side breaches.	Demonstrate active governance, HITL oversight, and validation; maintain own E&O / cyber insurance.

Domain	What DPA/MBSA Provides	What the Firm Must Still Do
Data Residency	Contractual assignment of risk under SCCs and EUDB commitments.	Technically enforce residency through admin routing, Multi-Geo, and CA controls; evidence configuration.
Audit Trails	UAL metadata logging and Service Trust Portal attestations.	Operate WORM export pipelines, configure Litigation Hold, retain prompt/response content for regulator review.
Model Safety	Microsoft's internal safety filters and content-safety scoring.	Independent model validation, output verification, drift monitoring, model inventory.
Incident Response	Microsoft SLA notification duties under the DPA.	Internal DORA 4h/24h/72h/1mo playbooks, NYDFS 72h notification, SIEM integration, legal/compliance escalation.
Sub-Processor Changes	Notice of new sub-processors with right to object.	TPRM intake on each notice, residual risk assessment, business-line concurrence, register update.
Recordkeeping	Native retention configuration tools.	Configure retention to firm's longest applicable horizon, evidence configuration, run quarterly retention attestations.

The bottom line is that contracts shift legal liability, but regulatory compliance requires demonstrable operational control. Until the firm actively configures retention pipelines, enforces residency boundaries, monitors DSPM alerts, validates model outputs, and exercises its incident response plan against Copilot- and CoWork-induced scenarios, the DPA and MBSA alone will not satisfy a G-SIFI regulatory audit.

Part 4 — Recommended Controls Catalog

The catalog below is a regulator-ready summary of the controls a G-SIFI deploying M365 E5 with Copilot (Anthropic) and CoWork should operate. Tier 1 controls are mandatory under one or more of the regulations cited; Tier 2 controls are strongly recommended industry practice; Tier 3 controls reflect mature-state operations. The firm's specific posture may add, or substitute controls based on its jurisdictional footprint.

4.1 Tier 1 — Mandatory

Control	Primary Authority	M365 Mechanism
Phishing-resistant MFA for all users	NYDFS §500.12	Conditional Access + FIDO2 / Windows Hello / certificate-based authentication.
Privileged access management with PIM and just-in-time activation	NYDFS §500.7(c) (Class A), SOX §404 ITGC	Entra PIM, Access Reviews.
Third-party service provider security policy and register	NYDFS §500.11, DORA Art. 28	TPRM intake, register including Anthropic and Microsoft sub-processors.
Asset inventory and data retention policy	NYDFS §500.13	M365 Admin Center, Records Management, asset register.
Incident response plan with 72h notification	NYDFS §500.16, §500.17	Defender XDR + Sentinel runbooks, legal escalation.
DORA 4h/24h/72h/1mo incident reporting	DORA Articles 17–23	Sentinel-driven classification + NCA reporting workflow.
Books-and-records preservation	SEC 17a-4, FINRA 4511	Records Management labels, Litigation Hold, WORM archive.
Communications recording	MiFID II Art. 16(7), FINRA 3110	Teams recording, Communication Compliance, retention.
Model risk management for AI use cases	SR 11-7, OCC 2011-12, EU AI Act	AI Risk Committee, model inventory, independent validation.
GDPR DPIA for Copilot/Anthropic in EU	GDPR Art. 35	Documented DPIA, residual-risk sign-off.

Control	Primary Authority	M365 Mechanism
EUDB enabled where applicable	GDPR Chapter V, EDPB	M365 Admin Center setting, evidenced configuration.

4.2 Tier 2 — Strongly Recommended

Control	Driver	M365 Mechanism
Sensitivity Labels mandatory across mail, files, and Teams	Industry practice; supports DLP, eDiscovery, IRM	Purview Information Protection.
DLP policies covering Copilot prompt and response paths	Industry practice; FFIEC, GLBA Safeguards	Purview DLP with Copilot endpoints enabled.
Customer Lockbox approval workflow	Operational resilience, regulator confidence	Customer Lockbox enabled with documented approver chain.
Insider Risk Management with departing-user signal	Industry practice	Purview Insider Risk.
Communication Compliance over Copilot content	FINRA 3110, supervisory expectations	Purview Communication Compliance.
DSPM-based Graph permission exposure review	NYDFS §500.7, industry practice	Purview DSPM.
Quarterly access reviews on Copilot and CoWork app permissions	SOX ITGC, NYDFS §500.7	Entra Access Reviews.
Document the Copilot Decommissioning Runbook and exercise annually	DORA exit strategy, FCA SYSC 8	Internal runbook + tabletop exercise.

4.3 Tier 3 — Mature-State

M365 E5 ARCHITECTURE WITH COPILOT OVERLAY – DRAFT

Control	Driver	M365 Mechanism
Customer Key for Microsoft 365 (BYOK)	Highest assurance posture; some sovereign tenants	Customer Key with HSM-backed Azure Key Vault.
Token Protection (binding) for high-risk apps	Defence in depth	Conditional Access Token Protection.
Threat-led penetration testing under TIBER-EU / CBEST scope	DORA Art. 26-27, BoE CBEST	Internal red team + external provider, with M365 attack-path coverage.
Bring-your-own model evaluation harness for Copilot use cases	Mature MRM	CoWork-hosted evaluation pipeline; output A/B testing.
Information Barriers for in-scope research/banking populations	Chinese walls, MAR	Purview Information Barriers.
Conditional Access App Control over M365 Admin Center settings	Configuration drift prevention	Defender for Cloud Apps + CA App Control.

Appendix A — Glossary

Term	Definition
Copilot CoWork	The firm's internal extension layer that surfaces curated agentic skills, brokers OBO tokens, and enforces firm-specific policy in front of or alongside M365 Copilot.
DPA	Microsoft Products and Services Data Protection Addendum — the contractual document that governs Microsoft's processing of Customer Data.
DPIA	Data Protection Impact Assessment under GDPR Article 35; required for processing likely to result in elevated risk to data subjects.
EUDB	EU Data Boundary — Microsoft's commitment to store and process Customer Data and pseudonymized personal data within the EU/EFTA for in-scope services.
Grounding	The retrieval step in which Copilot pulls context from the user's authorized Graph corpus to construct an inference payload.
MBSA	Microsoft Master Business and Services Agreement; the umbrella commercial agreement under which the DPA and Online Services Terms operate.
Multi-Geo	M365 capability allowing user data to be pinned to satellite regions outside the tenant's home region.
OBO	On-Behalf-Of OAuth 2.0 token exchange; allows a middle-tier service to call a downstream API as the originating user.
PIM	Privileged Identity Management; Entra capability for just-in-time, approval-gated activation of privileged roles.
PRT	Primary Refresh Token; device-bound refresh token issued by Entra ID for Windows clients.
Substrate	The Exchange Online-backed system mailbox layer that stores hidden, system-generated content including Copilot prompt/response artefacts.
TPRM	Third-Party Risk Management.
UAL	Unified Audit Log (now Microsoft Purview Audit).

Term	Definition
WAM	Web Account Manager; the Windows broker that integrates Entra authentication with the OS.
WORM	Write-Once-Read-Many; a storage mode in which records cannot be altered or deleted before their retention horizon.
EDP	Enterprise Data Protection; Microsoft’s contractual commitment under the DPA and Product Terms that Customer Data is not used to train models and is not retained beyond inference.

Appendix B — Regulatory Citation Index

Citation	Title / Subject
17 CFR 240.17a-4	SEC Rule 17a-4 — Records to be preserved by certain exchange members, brokers and dealers.
23 NYCRR Part 500	NYDFS Cybersecurity Requirements for Financial Services Companies (Second Amendment effective 1 November 2023).
§500.7	Access Privileges and Management.
§500.11	Third-Party Service Provider Security Policy.
§500.12	Multi-Factor Authentication.
§500.13	Asset Management and Data Retention Requirements.
§500.16	Incident Response Plan.
§500.17	Notices to Superintendent (72-hour notification).
DORA / Regulation (EU) 2022/2554	Digital Operational Resilience Act — applies from 17 January 2025.
DORA Articles 17–23	ICT-related incident management and reporting framework.
DORA Articles 28–30	ICT third-party risk; critical ICT third-party providers.
EU AI Act / Regulation (EU) 2024/1689	Harmonized rules on artificial intelligence.
FINRA Rule 4511	Books and Records — General Requirements.
FINRA Rule 3110	Supervision.
GDPR / Regulation (EU) 2016/679	General Data Protection Regulation; Chapter V governs international transfers.
MiFID II / Directive 2014/65/EU	Article 16(7) — recording of telephone and electronic communications.
OCC Bulletin 2011-12	Sound Practices for Model Risk Management — companion to SR 11-7.

Citation	Title / Subject
SR 11-7	Federal Reserve Supervisory Guidance on Model Risk Management.
SOX §404	Sarbanes–Oxley Act, Internal Controls over Financial Reporting.
MAS TRM Guidelines	Monetary Authority of Singapore — Technology Risk Management Guidelines.
HKMA SA-2	Hong Kong Monetary Authority — Outsourcing Supervisory Policy Manual.
FCA SYSC 8	UK Financial Conduct Authority — Outsourcing chapter of the Senior Management Arrangements, Systems and Controls sourcebook.

Appendix C — Sub-Processor Reference (M365 Copilot)

The table below reflects Microsoft's confirmed M365 Copilot sub-processor disclosures as of the baseline date. Microsoft's authoritative list lives in the DPA Sub-Processor List on the Microsoft Trust Center; the firm's TPRM team should reconcile this table against that list at each review cycle.

Sub-Processor	Role	Geographic Scope	Firm Posture
Microsoft (intra-company)	Operates Azure OpenAI; primary inference path.	Global, including EUDB.	Default; included in core DPA.
Anthropic	Provides Claude models for Copilot inference.	Disabled by default in EU/EFTA/UK; unavailable in government/sovereign clouds; default-on in US tenants with admin opt-out.	Enabled only after DPIA completion and TPRM concurrence; pinned through admin policy.
xAI	Provides Grok models. Independent processor — NOT under the Microsoft M365 Copilot DPA. Reachable via Copilot Studio integration only, not the M365 Copilot inference path.	US tenants only.	Treat as a firm direct vendor relationship: TPRM register entry, separate DPA/MSA, firm DPIA. Microsoft DPA flow-down does not apply.

Other model providers

Models such as Mistral, Meta Llama, DeepSeek, and Alibaba Qwen are present in the Azure AI Foundry catalog under their own Azure terms. They are not, today, confirmed M365 Copilot sub-processors with active DPA flow-down for M365 Customer Data. If the firm chooses to use any such model from CoWork or another firm-operated workload, that relationship is governed under separate Azure (or partner) terms and is the firm's vendor relationship, not Microsoft's.

Appendix D — Architecture-to-Obligation Cross-Reference

Architecture Element	Tier	Primary Regulatory Anchors
Microsoft Entra ID + Conditional Access + PIM	Identity plane	NYDFS §500.7, §500.12; SOX §404 ITGC; ISO 27001 A.9; FFIEC Authentication.
Exchange Online / SharePoint / OneDrive / Teams (incl. Substrate)	Data plane	GDPR; SEC 17a-4(b)(4); FINRA 4511; MiFID II Art. 16(7); state breach laws.
Microsoft Purview (Sensitivity, DLP, Records, eDiscovery, CC, Insider Risk, DSPM)	Compliance plane	GDPR; NYDFS §500.13; SEC 17a-4; FINRA 3110/4511; SR 11-7 (where compliance-decision data is in scope).
Purview Audit + Defender XDR + Entra logs	Telemetry plane	DORA Art. 12, 17–23; NYDFS §500.16, §500.17; SOX ITGC; SR 11-7 (audit evidence).
EUDB + Multi-Geo + Customer Lockbox	Residency	GDPR Chapter V; EDPB Schrems II; DORA Art. 28(7); MAS TRM 11; HKMA SA-2.
Copilot orchestrator	AI overlay	EU AI Act Art. 26; SR 11-7; OCC 2011-12; sectoral AI guidance.
Anthropic sub-processor path	AI overlay	GDPR; NYDFS §500.11; DORA Art. 28-30; firm DPIA.
Copilot CoWork	Firm-operated	Full firm responsibility — SR 11-7, EU AI Act, GDPR, NYDFS §500.7/11/13/16, SEC 17a-4, FINRA 4511.

Appendix E — Document Acceptance

By acknowledging this document, the named function confirms that it has reviewed the architecture description and the regulatory mapping for accuracy in respect of its area, and that any material divergence from the firm's existing posture has been raised through the standard governance forum.

Function	Reviewer Name	Date	Sign-off
Enterprise Architecture			
Chief Information Security Officer			
Privacy Office			
Model Risk Management			
Third-Party Risk Management			
Records & eDiscovery			
Operational Resilience / DORA Lead			
Compliance — Communications Surveillance			

End of document.